Data & Democracy

Political microtargeting: A threat to electoral integrity?

Dobber, T.

Publication date
2020

Document Version
Other version

License
Other

Chapter 6

Conclusion

Key findings

Against the backdrop of surprising electoral outcomes in the UK and the US in 2016, this dissertation set out to learn more about PMT: an opaque technique, very American, and in some popular accounts hyped as 'weaponized ad technology' (Singer, 2018) or digital mindreading (Soares, 2016). While such hyperboles have no foot in reality, there are legitimate concerns about PMT's potential pitfalls that are worth the scientific scrutiny. In the coming section, I describe four main problems in the field of microtargeting research, and briefly explain how my key findings each inform the problems.

One of the larger problems in the field of microtargeting research, from the outset, was a complete blind spot about what PMT would look like in a non-Anglo-Saxon, multi-party, small-budgeted (i.e., Dutch) context. Some scholars explicitly wondered whether political parties can "campaign in Europe as they do in North-America" (Bennett, 2016; p. 1). Chapter two proposed and tested a model that can be used to understand how contextual factors hinder or facilitate data-driven capabilities of political campaigns in different contexts.

> *Key finding 1: Dutch parties can and do microtarget, but contextual factors mold PMT into a Dutch variant.*

Dutch parties operate in a proportional representation, multi-party context. While there are no 'battleground states' like in the US, Dutch political parties do identify 'key areas' that require relatively intensive campaigning efforts. These are areas where, for example, turnout is low, but voting potential is high. Data-analysis helps identifying these areas and microtargeting techniques help reach specific citizen groups.

While Dutch parties have little money, in comparison with their German or Belgian counterparts, party budgets do not need to restrain the microtargeting efforts of political parties, but the small budgets do shape them. Where one large Dutch party refrained from adopting microtargeting capabilities beyond basic Facebook-based targeting because of their relatively small budget, a different party made the best out of its strapped-for-cash situation by organizing its campaign in a citizen-initiated manner (see Gibson, 2015). This resulted in bottom-up creativity that engaged supporters on the one hand and could be used as content to microtarget to potential voters on the other hand. In sum, having a small budget does not make microtargeting impossible. Instead, it can facilitate a context-dependent form of microtargeting.

A second problem in the field of microtargeting research was that there were some normative and theoretical claims about how large-scale data-collection and subsequent communicating of tailored messages could cause a perceived loss of intellectual privacy. People would experience the disappearance of a room to think for oneself if they grew

concerned about their privacy. But this was never empirically tested in a European setting.

> *Key finding 2: There is a downward spiral, where changes in privacy concerns lead to changes in attitudes toward PMT and vice versa, over time.*

This downward spiral is the first evidence for more normative and theoretical claims about a perceived loss of intellectual privacy in relation to PMT (see chapter 3). This spiral is important because it shows us how citizens could be starting to behave in an undesirable way from a democratic perspective. This downward spiral could, for instance, result in chilling effects or hinder people in thinking independently and deviate from the popular opinion.

A third problem in the microtargeting literature was that the effectiveness of PMT techniques were assumed in the literature (e.g. Gorton, 2016) and the media (Cadwalladr, 2017), but never measured. While, in 2019, US-based research found evidence for effectiveness of PMT-techniques (Haenschen & Jennings, 2019; Endres, 2019), there still was no empirical knowledge about PMT's effectiveness in a non-US, multiparty context.

> *Key finding 3: PMT can be reasonably effective in increasing the likelihood to vote for the political party sending the microtargeted message.*

Recipients of the PMT message on average increased ten percent in likelihood to vote for the microtargeting party, in comparison with the likelihood to vote of the control group. Chapter 4 describes a field experiment that took place in Utrecht, during the municipality elections of 2018. The chapter offers a first glimpse into PMT's effects in a natural multi-party election context.

A fourth problem was that a new understudied disinformation mode rose through the ranks as a potentially very manipulative force: deepfakes. Especially manipulative when combined with microtargeting. Political actors could expose people to deepfakes in which they could try to damage a politician. Microtargeting techniques could amplify the damage by allowing malicious actors to send different, tailored deepfakes to only those people who are likely susceptible to the specific deepfake.

> *Key finding 4: Microtargeting techniques can further amplify the negative effects of a deepfake.*

Using a self-produced deepfake of a Dutch politician, I ran an online experiment where I measured the effects of exposure to a (microtargeted) deepfake on political attitudes (see chapter 5). The experiment showed that people exposed to the deepfake held significantly worse attitudes, than the control group. Moreover, microtargeting the deepfake amplified its effects on political attitudes, but only for a relatively small group of participants. For now it seems that microtargeted deepfakes *can* become a problem in the near future. But the complexity of an election makes it challenging to use microtargeted deepfakes to achieve specific goals in a controlled way. Getting accurate perceptions of who is susceptible to what information is vital, but not easy to do.

A threat to electoral integrity?

PMT can potentially be used to improve the integrity of elections. It can help convey political ideas. It can potentially help reach and activate hard-to-reach groups. It may help improve political knowledge. The positive potential is not the focus of this dissertation, but it should be studied because that would help us understand how PMT can be a force for good. This dissertation focuses on five threats to electoral integrity: privacy, intellectual privacy, manipulation, fragmentation of the public agenda and changing political power structures. The first and second threat ([intellectual] privacy) are 'gateway threats' that open the door to the other threats. The threat of manipulation follows directly from these and applies to the individual citizen. Fragmentation of the public agenda is a process that is much less focused on the individual citizen. The changing of political power structures occurs on the most abstract level. In the coming sections I will discuss how the research in this dissertation has informed our understanding of these five threats to electoral integrity.

Privacy

The General Data Protection Regulation (GDPR) was adopted in 2016, but came into force in 2018. The GDPR is meant to ensure transparency and fairness in data collection and processing. But chapter two of this dissertation shows how some of the practices of political parties are questionable from GDPR point of view. For example political parties, rather than collecting personal data themselves, can rely on the services by social media platforms, such as Facebook. Facebook but also Google require an *omnibus* consent of its users, but it is not clear whether that consent qualifies as *informed* consent as specified in the GDPR. More importantly, social platforms offer political advertisers the possibility to target many different characteristics that strongly correspond with someone's political preferences without necessarily collecting special categories of data in the sense of article 9 GDPR (see Wachter, 2020; Youyou, Kosinski & Stillwell, 2015). This raises the question to what extent the ban on collecting special categories of data (including information about political preferences) is still effective in protecting people from discriminatory or otherwise invasive practices in political microtargeting. Even by not explicitly targeting someone because they sympathize with the Green Party, but instead targeting someone because they regularly shop at a vegan market and donate to

Greenpeace, political parties can use a proxy of political preference without obtaining specific informed consent from the targeted citizen (as required by article 9 GDPR).

Even if political parties want to adhere to the GDPR, and perhaps in the near future the revised e-Privacy Regulation, parties often lack the legal knowledge to adhere to the GDPR. As one campaign leader puts it in chapter two: "Legislation has grown so very comprehensive and complex. It's almost impossible to cope for us as a small organization". Political parties want to be compliant with regulation meant to safeguard the right to data protection, but are not (always) equipped to do so. Such a compliance deficit is problematic as it can infringe upon at least two of the Venice criteria on good practices in elections: free and secret suffrage (Council of Europe's Venice Committee; 2002). Political parties need more guidance on this issue (see also the call for guidance by the UK's Information Commissioner's Office, 2018).

Privacy infringement threatens the integrity of elections, as it can violate the practices of free and secret suffrage. Privacy infringement is especially threatening to the integrity of elections when the actor is operating in bad faith Actors who do not want to comply with regulation and who have anti-democratic goals are malicious actors. These can be foreign states, or (domestic) anonymous interest groups. Malicious actors can leverage (special categories of) personal data to amplify the already damaging effects of disinformation. When microtargeting disinformation, such as a deepfake, personal data are used in an unfair way: with the goal of deceiving citizens and furthering one's opaque political goals. This shows the importance as well as the limits of law in a connected world. Data protection law in general protects voters against the collection, profiling and microtargeting practices that do not comply with the key principles of data protection law (legitimate ground, purpose specification, data minimalization, protection of special categories of data, privacy by default, and rights of the data subjects). But data protection law does not protect against the use of data to unfairly manipulate voters. This is a problem with regards to the integrity of elections vis-à-vis platforms and political parties. It is even more problematic in the context of malicious actors who use PMT for anti-democratic purposes.

Intellectual privacy

Intellectual privacy is the perception of having a room to think for oneself (Dawes, 2014). It entails the freedom people perceive to deviate from dominant opinion. Chapter 3 shows that citizens hold unfavorable attitudes toward PMT and these attitudes deteriorate over time, while reciprocally privacy concerns grow. These negative citizen perceptions can be threatening to the integrity of elections because they form a downward spiral that can result in democratically undesirable outcomes such as chilling effects. The downward spiral found in chapter 3 points to the disappearance of one's room to think for oneself (see Dawes, 2014), which may decrease citizens' feeling of autonomy and their ability to independently form non-common opinions. Autonomous

citizens are a cornerstone of liberal democracy, and, therefore, as Gavison (1980, p. 455) puts it "to the extent that privacy is important for autonomy, it is important for democracy as well."

It could be that media reports about PMT that claim digital "mind readers" (Soares, 2016) "know you better than you know yourself" (Brennan, 2012), have induced a feeling of domination within citizens (Pettit, 2018). Or the feeling that their fellow citizen is easily fooled by bad faith PMT. Karpf (2019) argued that elite perceptions of digital techniques may perversely lead to short-term focused, self-interested, nihilistic behavior by politicians. On the basis of chapter 3, I would advise future research to extend Karpf's politician-focused argument to citizens. Citizen perceptions of PMT can potentially be threatening to the integrity of elections, because they can make them doubt the decision of their fellow citizens: why would someone respect the vote of their fellow citizens if those fellow citizens are so easily fooled by PMT? What is an election result worth if citizens believe elections are not decided in favor of the politicians with the best ideas, best past performances, or even the most charismatic leader, but rather by who has the most digitally advanced techniques? It reminds me of the *Tour de France*. What does it mean to carry the yellow jersey to the finish line in Paris, when all cyclists are perceived to use doping? And, more importantly, would people respect the winner? Or pay attention to next year's *Tour de France*? More research on the perceptions of PMT's effectiveness and the effects of those perceptions on electoral integrity is needed.

Manipulation

We know that communication effects are generally small (Bennett & Iyengar, 2008; Kalla & Broockman, 2018). People come into contact with different information in different ways. It is difficult to attend to all information we are exposed to. We also know that communication effects are often conditional (Valkenburg, Peter, & Walther, 2016): effects almost never apply equally to the population. PMT techniques can relatively successfully leverage these conditionalities so that subsegments of people are affected to a larger extent by a political message. PMT makes political advertising not only more efficient (see chapter 2), but also more effective (chapters 4 and 5). Making the dissemination of political information more efficient and effective is not inherently good or bad, but tools that give political actors more punching power need to be scrutinized closely upon their potential to impact electoral integrity.

PMT messages are by no means a hypodermic needle. But chapter four and five show that PMT's effects are not minimal either (that is, *if* the right message indeed reached the right person, something that is not easy to pull off). PMT alone will never persuade a rightist VVD voter to switch to the leftist PvdA. But PMT can be a tool to influence voters within their consideration sets (see Rekker & Rosema, 2019). This can be something positive when used in good faith. Legitimate actors should be able to inform

or even persuade citizens with the reasonable effectiveness offered by PMT techniques (see chapter 4).

But PMT's effectiveness can easily become manipulative when used in bad faith. For instance to spread deepfakes, lies, or misleading messages. Manipulation is "the use of information technology to covertly influence another person's decision-making, by targeting and exploiting their decision-making vulnerabilities" (Susser, Roessler & Nissenbaum, 2019, p. 1).

Using PMT in bad faith means leveraging people's personal information to influence them on false grounds, in pursuit of anti-democratic goals. If PMT were ineffective, quod non, bad faith PMT would only be problematic because the bad faith actor pursues anti-democratic goals. But people's political behavior and attitudes would remain unaffected. However, this dissertation shows PMT can be reasonably effective, and therefore vulnerable to misuse. Moreover, it seems as if PMT's effectiveness can influence different kind of socio-demographic groups. For instance, chapter 4 shows that a rather highly educated sample was affected by PMT. Chapter 5 shows that only a handful of participants recognized the deepfake as being manipulated. Hence, PMT's effectiveness can be manipulative for a large part of the electorate and thus threaten electoral integrity, if the technique is used in bad faith.

Fragmentation of the public agenda

Dutch statesman Johan Thorbecke meant for political parties in the Netherlands to campaign for the general interest rather than local interests, and cater to heterogeneous groups rather than homogeneous groups. But PMT is the exact opposite of what Thorbecke had in mind when he wrote the Dutch constitution and electoral law halfway the 1800's (Kaal, 2016). In that sense, in theory, aligning messages to small homogeneous subgroups of society is a threat to the integrity of the Dutch elections. In practice, the research in this dissertation finds no evidence for any fragmentation of the public agenda as a result of PMT. However, fragmentation is something that should be measured over a longer time-span, and on a larger scale than in chapter 4, which focused on four similar neighborhoods in Utrecht in the run-up to one municipality election. Moreover, microtargeting did not occur on a large scale in the Netherlands before 2016. While I found no evidence of fragmentation of the public agenda, I cannot with certainty conclude that there is or is not a threat to electoral integrity. Future research should avoid the limitations of chapter 4 and study fragmentation on a larger scale, for a longer time period.

Political power structures

In the run-up to the European parliamentary election of 2019, EU's top officials sent a letter to Facebook, urging the company to loosen their political advertising policy

to allow for pan-European advertising. The EU-officials hailed Facebook ads as "an important part of [its] digital information toolbox" (Welle, Tranholm-Mikkelsen & Selmayr, 2019). These developments suggest that the perceived usefulness of PMT holds not just for the Netherlands, but even for the EU itself. This brings up questions about the changing political power structures. Below, I discuss three sub-issues of concern. First, the commercialization of politics. Second, the political power of social platforms. Third, the political power of malicious actors.

The issue of commercialization

Commercialization of politics means perceiving citizens' votes as something that can be bought by political advertisers, and sold and profited from by data brokers and social platforms. Commercialization (or profit-maximization) drives the privacy-infringing data-collection and analysis industry, who are less concerned with electoral integrity than with shareholder value. Social platforms have been found to enable advertisers to tailor and target Jew haters (Angwin, Varner & Tobin, 2017), were slow to respond to signs of misuse for disinformation campaigns (Frenkel et al., 2018), and played a dubious role in facilitating communication in the run-up to genocide in Myanmar (Stevenson, 2018). Commercialization favors the political actors with deep pockets, potentially drowning out minority voices. In 2017, small-budgeted parties could still run a creative campaign and be heard (see chapter 2). But further commercialization makes that increasingly difficult. Moreover, commercialization further reduces citizens to targets (see chapter 3). And commercialization of politics is, in part, maintained by political advertisers, who perceive PMT as a useful technique (see chapter 2, but also 4, and 5). This perception leads to revenue for social platforms, further commercialization, and gives political power to social platforms.

The issue of political power of social platforms

Elite-perceptions of PMT's usefulness (see chapter 2) put social platforms in a position of dominance, where they can arbitrarily interfere in election campaigns: e.g., by charging more to certain political parties, or to reach specific voter groups (see Ali et al., 2019). Elite-perceptions of PMT's usefulness also potentially facilitate unchecked growth, and stymie competition. Why would the EU take action against an "important part of [its] digital toolbox" (Welle, Tranholm-Mikkelsen & Selmayr, 2019)? And why would the EU use the services of a competitor, when the current platform is already seen as useful? As a consequence, political actors advertise on only a few, ever expanding, ever powerful, platforms (see also Guess et al., 2019).

The political power of malicious actors

PMT helps malicious actors to further their political goals. Those political goals are being pursued by influencing people on the basis of lies, deceit and manipulation. Chapter 5 shows that political attitudes can be affected by deepfake disinformation and

PMT-techniques can further amplify those effects. These techniques together can, for instance, help malicious actors damage a politician who is critical toward their interests. Having the capability to create deepfakes, and to then microtarget those deepfakes, enables malicious actors to manipulate the target electorate into doing or thinking something in line with the interests of the malicious actor. Consequently, the malicious actor gains political power, essentially changing the power structures in politics by threatening the integrity of elections.

Future research

Throughout this section, I've provided three directions for future research. One, the positive potential of PMT. Two, research on the perceptions of PMT's effectiveness and the effects of those perceptions on electoral integrity. Three, study fragmentation on a larger scale, for a longer time period. But there are some other interesting perspectives on how electoral integrity can be maintained that deserve some scholarly follow-up.

Bad faith microtargeting

Effectiveness of political microtargeting is in itself not problematic, but effectiveness of bad faith microtargeting is. A potential research focus would be to examine the feasibility and desirability of a system where unverified or illegitimate actors are disallowed to politically microtarget, use lookalike audiences, and custom audiences functions. Although we have already seen that legitimate actors can use microtargeting in bad faith (e.g., by microtargeting fake advertisements), arguably, this measure would cut out many potential malicious actors and make the field easier to scrutinize. The philosophical question about what makes an actor (il)legitimate and what constitutes as bad faith microtargeting is crucial. Canada has already moved forward by prohibiting foreign actors to pay for political ads targeted at Canadian citizens, which makes Canada an interesting context to study the effectiveness of this measure (Elections Modernization Act, 2018). The Dutch government is preparing a similar measure: to prohibit non-EU gifts, but only gifts to political parties (Rijksoverheid, 2020).

Transparency

An important research question would be whether threats to electoral integrity can partly be mitigated by providing users with a concise and comprehensible report. Such a report should detail the tailored ads people have been exposed to themselves, and have clicked on; but also the ads their (close) friends saw and their click-through-rates as an aggregated number. This report should be shown on a regular basis in the month leading up to election date, in the style of Spotify wrapped, or Facebook's 'Year in review'. Likely, citizens would see that the click-through-rate is very low for themselves but also for their friends. Citizens would also better understand that their friends do not necessarily see the same ads as they themselves do.

Potentially, future research gives us more capacity to mitigate threatening applications of PMT, while utilizing PMT's positive potential. There is no easy fix that would solve all challenges. I am worried about the power of privately-owned, mostly American, companies that take on an ever-crucial role in political communication. I am also worried about the power of bad faith actors. I have my doubts about PMT's potential for manipulation on a grand scale, but I am more worried about small-scale manipulation that occurs under the radar. Finally, it is discomforting to know that new challenges will keep on arising alongside new technologies. But I am reassured by Cicero's observation that was true in 62 BC, is true now, and will be true in the future:

> *Nothing is more fickle than people in a crowd, nothing harder to discover than how men intend to vote, nothing trickier than the whole way in which elections work. (Cicero, Pro Murena, 62 B.C.)*